ISSN: 2993-5377

Research Article

Open Access Journal of Applied Science and Technology

Structural Isomorphism Between LLM Embedding Spaces and Quantum Mechanical Systems

Timo Aukusti Laine*

Financial Physics Lab, Finland

***Corresponding Author**
Timo Aukusti Laine, Financial Physics Lab, Finland.

Submitted: 2026, Apr 01; **Accepted**: 2026, May 14; **Published**: 2026, May 18

Citation: Laine, T. A. (2026). Structural Isomorphism Between LLM Embedding Spaces and Quantum Mechanical Systems. *OA J Applied Sci Technol, 4*(2), 01-18.

Abstract

Large language models (LLMs) represent semantic information as high-dimensional embedding vectors, typically compared using cosine similarity. This paper establishes a structural isomorphism between LLM embedding spaces and quantum mechanical systems. We demonstrate that the transformed cosine similarity can be precisely mapped to a quantum mechanical measurement probability, governed by a rank-1 Hamiltonian and the Schrödinger equation, where the time parameter is identified as a gauge variable. A commutative diagram unifies four equivalent expressions for this similarity, spanning classical bilinear forms and quantum Born rule probabilities. Crucially, the quantum formulation reveals a local U(1) gauge symmetry, absent from classical descriptions, with an associated conserved semantic charge modeling contextual influence. A quantum circuit on logarithmically many qubits yields this similarity as a single measurement probability, a result confirmed by simulation. The isomorphism is structural, not physical: it highlights that the mathematical structures governing LLM embeddings are precisely those of quantum mechanics, and this correspondence is exact.

1. Introduction

Large language models (LLMs) encode semantic information as high-dimensional embedding vectors, with cosine similarity serving as the standard measure of semantic relatedness. Despite their empirical power, LLM embedding spaces currently lack a rigorous theoretical framework comparable to, for instance, statistical mechanics for thermodynamics. Embeddings are often treated as points in $\mathbb{R}^N$ with no deeper mathematical structure beyond the dot product. This paper addresses this gap by establishing a precise structural isomorphism between LLM embedding spaces and quantum mechanical systems.

The primary goals of this work are multifaceted. We first provide a complete and rigorous derivation of both the LLM embedding classical system and its quantum mechanical counterpart, detailing aspects previously presented in abbreviated form. Crucially, we offer a clear interpretation of the quantum mechanical time parameter and complex phases within this context, resolving ambiguities from prior work. Central to our contribution, we prove the exact structural isomorphism between these classical and quantum formulations, demonstrating their mathematical equivalence in describing semantic similarity. This isomorphism not only unifies these descriptions but also reveals new mathematical structures, such as a local U(1) gauge symmetry, genuinely absent from classical accounts, and clarifies the profound implications of this correspondence. We also present a concrete quantum circuit implementation that leverages this isomorphism, providing an experimental verification of the theoretical framework by showing that classical cosine similarity can be precisely recovered as a quantum measurement probability. This isomorphism is structural, not physical: it highlights that the underlying mathematical principles governing LLM embeddings—vector spaces, inner products, projections, and squared norms—are precisely those of quantum mechanics, and that this correspondence is exact.

The paper is organized as follows. Section II reviews background and related work. Section III derives the LLM embedding classical system from the geometry of the unit sphere. Section IV constructs the quantum system and proves its equivalence to the classical system through the Schrödinger equation. Section V establishes the commutative diagram and characterizes the structural isomorphism. Section

VI resolves the interpretation of time as a gauge parameter and explains the role of the complex state vector. Section VII presents the quantum circuit implementation and its derivation from the Hamiltonian structure. Section VIII reports the experimental validation through quantum circuit simulations. Section IX discusses the implications of the framework, the additional quantum structure it provides, and directions for future work.

2. Background and Related Work

The mathematical framework developed in this paper draws on three distinct lines of research: the geometry of embedding spaces in natural language processing, the application of quantum formalism to meaning and cognition, and the practical intersection of quantum computing with machine learning.

The idea that meaning can be represented geometrically dates to the distributional hypothesis—the observation that words appearing in similar contexts tend to have similar meanings [1,2]. This hypothesis underlies the construction of word embedding spaces, where semantic relationships are encoded as geometric relationships between vectors. The development of contextual embeddings through transformer architectures and sentence-level representations extended this geometric picture to richer linguistic units, with cosine similarity between embedding vectors serving as the standard measure of semantic relatedness [3-6]. The present work takes this geometric structure as its starting point and shows that it admits a reformulation in the language of Hamiltonian mechanics and, subsequently, quantum mechanics.

The application of quantum-mechanical formalism to language and cognition has a separate history. Widdows [7] explored connections between vector space models of meaning and quantum logic, exploiting the shared mathematical structure of subspaces, projections, and orthogonality. Coecke, Sadrzadeh, and Clark [8] developed a compositional distributional model of meaning grounded in category theory, using the tensor product structure of quantum mechanics to compose word meanings into sentence meanings. In the cognitive sciences, quantum probability theory has been applied to model phenomena in human judgment and decision-making that violate classical probability axioms, such as order effects in sequential judgments and the conjunction fallacy. These approaches share with the present work the recognition that quantum-mechanical mathematics can describe systems that are not themselves quantum physical, but they differ in scope and method: they typically postulate quantum-like structures to model specific linguistic or cognitive phenomena, whereas the present work derives the quantum structure from the existing geometry of LLM embeddings without additional postulates.

On the computational side, quantum computing offers a distinct set of primitives—superposition, entanglement, and interference—that have been applied to machine learning tasks including classification, clustering, and kernel methods [9,10]. Quantum algorithms for search and optimization demonstrate that quantum resources can provide advantages for specific computational problems [11,12]. The quantum circuit constructed in this paper belongs to this computational tradition: it encodes classical embedding data into quantum states and extracts cosine similarity as a measurement probability. However, the circuit is derived from the structural isomorphism rather than designed as a standalone algorithm, and its primary purpose is to provide an experimental verification of the theoretical framework rather than a computational speedup.

The present work builds directly on a series of earlier papers by the author [13–19]. In Ref. [13], the concept of a semantic wave function was introduced, associating quantum state vectors with LLM embeddings. Ref. [14] developed the quantum LLM framework, modeling semantic spaces with quantum principles. Ref. [15] demonstrated the first quantum circuit implementation on IBM quantum hardware, computing cosine similarity as a measurement probability on a 128-qubit processor. Ref. [16] introduced the Hamiltonian formulation and the discrete semantic state structure that forms the classical foundation of the present work. Refs. [17] and [18] extended the framework to partition function similarity and hierarchical models of linear and nonlinear dynamics in embedding spaces, and a local U(1) gauge symmetry. Ref [19] shows how to implement quantum algorithms for Large Language Models on noisy intermediate Scale quantum computers. The present paper consolidates and extends these results by providing complete derivations of both the classical and quantum systems, proving their structural isomorphism through a commutative diagram, resolving the interpretation of time and complex phases, and presenting a systematic experimental validation, which were previously given in abbreviated form.

3. The LLM Embedding Classical System

In this section we derive the classical framework that forms the foundation of the quantum mapping. Starting from the standard definition of cosine similarity between two L2-normalized embedding vectors, we arrive at a Hamiltonian formulation with a universal eigenstructure through a direct geometric construction. The key result is that the transformed cosine similarity S'_C equals the square of the projection of $\mathbf{a}$ onto the angle bisector of $\mathbf{a}$ and $\mathbf{b}$, which we write as $S'_C = a'^2_1$. This form admits a direct and exact quantum mechanical interpretation, as we show in Section IV.

3.1. Cosine Similarity and the Embedding Vector Parameterization

We consider two embedding vectors $\mathbf{a}$ and $\mathbf{b}$ in $\mathbb{R}^N$, where N is the embedding dimension. The cosine similarity between $\mathbf{a}$ and $\mathbf{b}$ is defined as

$$S_C(\mathbf{a}, \mathbf{b}) = \cos\theta = \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \, \|\mathbf{b}\|} \tag{1}$$

We parameterize the two embedding vectors as

$$\mathbf{a} = [a_1, a_2, \ldots, a_N]^T \tag{2}$$

$$\mathbf{b} = [-a_1 + v_1\Delta, \; -a_2 + v_2\Delta, \; \ldots, \; -a_N + v_N\Delta]^T \tag{3}$$

where $\Delta > 0$ is a positive scalar and v_i are real coefficients satisfying $|v_i| \geq 1$ or $v_i = 0$. This parameterization is fully general: any two embedding vectors can be expressed in this form by appropriate choice of $\mathbf{v} = [v_1, v_2, \ldots, v_N]^T$ and Δ. The vector **b** is written as a perturbation from the antipodal point –**a**, with the perturbation direction and magnitude encoded in **v** and Δ, respectively.

3.2. The Fundamental Scale Δ and L2 Normalization

We assume that both vectors are L2 normalized

$$\|\mathbf{a}\|^2 = \sum_{i=1}^{N} a_i^2 = 1 \tag{4}$$

$$\|\mathbf{b}\|^2 = \sum_{i=1}^{N} (-a_i + v_i\Delta)^2 = \sum_{i=1}^{N} \left(a_i^2 - 2a_i v_i \Delta + v_i^2 \Delta^2\right) = 1 \tag{5}$$

Subtracting Eq. (4) from Eq. (5) yields the normalization condition

$$\sum_{i=1}^{N} v_i \left(-2a_i + v_i\Delta\right) = 0 \tag{6}$$

The parameter Δ plays an important role in this derivation. It is not an arbitrary or infinitesimally small quantity; rather, it reflects the inherent discrete structure of the token representations in the embedding space. Each token occupies a finite region in the embedding space, and L2 normalization rescales these regions onto the unit sphere. For practical calculations, Δ can be chosen as the smallest absolute value among the nonzero components of **a** and **b**.

3.3. Geometric Derivation of S'_C

The transformed cosine similarity S'_C can be derived through a direct geometric construction, as shown here, without explicit dependence on Δ. An alternative derivation, which incorporates Δ, is presented in Section IIIE. Define the unit vector along the angle bisector of **a** and **b**

$$\hat{\mathbf{v}} = \frac{\mathbf{a} + \mathbf{b}}{\|\mathbf{a} + \mathbf{b}\|} \tag{7}$$

This is well-defined whenever **a** + **b** ≠ 0, i.e., whenever **a** ≠ −**b**. The angle bisector is the unique direction that is equidistant from **a** and **b** on the unit sphere, making it the natural geometric object encoding the relationship between the two embedding vectors. The squared norm of **a** + **b** is

$$\|\mathbf{a} + \mathbf{b}\|^2 = \|\mathbf{a}\|^2 + 2\mathbf{a} \cdot \mathbf{b} + \|\mathbf{b}\|^2 = 2(1 + \mathbf{a} \cdot \mathbf{b}) \tag{8}$$

The projection of a onto $\hat{\mathbf{v}}$ is

$$\hat{\mathbf{v}} \cdot \mathbf{a} = \frac{(\mathbf{a} + \mathbf{b}) \cdot \mathbf{a}}{\|\mathbf{a} + \mathbf{b}\|} = \frac{1 + \mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a} + \mathbf{b}\|} \tag{9}$$

where we used (**a** + **b**) • **a** = ‖**a**‖² + **a** • **b** = 1 + **a** • **b**. Squaring Eq. (9) and substituting Eq. (8)

$$(\hat{\mathbf{v}} \cdot \mathbf{a})^2 = \frac{(1 + \mathbf{a} \cdot \mathbf{b})^2}{2(1 + \mathbf{a} \cdot \mathbf{b})} = \frac{1 + \mathbf{a} \cdot \mathbf{b}}{2} = S'_C \tag{10}$$

The transformed cosine similarity equals the squared projection of **a** onto the angle bisector of **a** and **b**. This is a direct consequence of the geometry of the unit sphere in $\mathbb{R}^N$ and requires no additional assumptions beyond L2 normalization.

Note that S'_C is symmetric in **a** and **b**: since $\hat{\mathbf{v}}$ is the bisector of both vectors, an identical calculation with **a** and **b** interchanged gives

$$\hat{\mathbf{v}} \cdot \mathbf{b} = \frac{(\mathbf{a}+\mathbf{b}) \cdot \mathbf{b}}{\|\mathbf{a}+\mathbf{b}\|} = \frac{1+\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}+\mathbf{b}\|} \tag{11}$$

which is identical to Eq. (9), so that $(\hat{\mathbf{v}} \cdot \mathbf{b})^2 = S'_C$ as well. The quantum circuit implementation of Section VII encodes **a** and **b** asymmetrically—**a** as the quantum state and **b** through $\hat{\mathbf{v}}$—but this symmetry guarantees that the measurement probability is independent of which vector is chosen as the initial quantum state.

3.4. The Rank-1 Hamiltonian and Its Eigenstructure

We cast the result of Eq. (10) into a Hamiltonian form

$$\mathbf{H}' = \hat{\mathbf{v}}\hat{\mathbf{v}}^T \tag{12}$$

This is the orthogonal projection matrix onto the one-dimensional subspace spanned by $\hat{\mathbf{v}}$. This Hamiltonian is related to the original cosine similarity Hamiltonian **H** (defined by **b** = **Ha**) through the linear shift

$$\mathbf{H}' = \frac{1}{2}(\mathbf{H}+\mathbf{I}) \tag{13}$$

where **I** is the $N \times N$ identity matrix. By construction, using the orthogonality $\mathbf{U}^T = \mathbf{U}^\dagger$,

$$\mathbf{a}^T\mathbf{H}'\mathbf{a} = \mathbf{a}^T(\hat{\mathbf{v}}\hat{\mathbf{v}}^T)\mathbf{a} = (\hat{\mathbf{v}} \cdot \mathbf{a})^2 = S'_C \tag{14}$$

so that S'_C is a quadratic form in a with matrix $\mathbf{H}'$

$$S'_C = \begin{bmatrix} a_1 & a_2 & \cdots & a_N \end{bmatrix} \frac{1}{\sum_{i=1}^N v_i^2} \begin{bmatrix} v_1^2 & v_1v_2 & \cdots & v_1v_N \\ v_2v_1 & v_2^2 & \cdots & v_2v_N \\ \vdots & \vdots & \ddots & \vdots \\ v_Nv_1 & v_Nv_2 & \cdots & v_N^2 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_N \end{bmatrix} \tag{15}$$

The eigenvalue structure of $\mathbf{H}'$ follows immediately from its rank-1 outer product form:
- $\mathbf{H}'$ has rank 1, so it has exactly one nonzero eigenvalue.
- The trace satisfies $\mathrm{Tr}(\mathbf{H}') = \hat{\mathbf{v}}^T\hat{\mathbf{v}} = \|\hat{\mathbf{v}}\|^2 = 1$, so the nonzero eigenvalue equals 1.
- The remaining $N-1$ eigenvalues are all zero.

The eigenvector corresponding to eigenvalue 1 is $\hat{\mathbf{v}}$ itself, since $\mathbf{H}'\hat{\mathbf{v}} = \hat{\mathbf{v}}(\hat{\mathbf{v}}^T\hat{\mathbf{v}}) = \hat{\mathbf{v}}$. The $N-1$ eigenvectors corresponding to eigenvalue 0 span the orthogonal complement of $\hat{\mathbf{v}}$. This eigenvalue structure {1,0,0,...,0} is universal: it holds for any pair of embedding vectors, with all pair-specific information encoded in the eigenvectors rather than the eigenvalues.

3.5. Derivation of S'_C with Δ

Using the L2 normalization conditions, the cosine similarity between the normalized vectors simplifies to

$$S_C = \mathbf{a} \cdot \mathbf{b} = \sum_{i=1}^N a_i(-a_i + v_i\Delta) = -\sum_{i=1}^N a_i^2 + \Delta\sum_{i=1}^N a_iv_i = -1 + \Delta\sum_{i=1}^N a_iv_i \tag{16}$$

Applying the normalization condition Eq. (6), we obtain

$$S_C = -1 + \frac{\Delta^2}{2}\sum_{i=1}^N v_i^2 \tag{17}$$

Since cosine similarity takes values in [−1, 1], we define a rescaled quantity

$$S'_C = \frac{1}{2}(S_C + 1) = \frac{\Delta^2}{4}\sum_{i=1}^{N} v_i^2 \tag{18}$$

which takes values in [0, 1]. Using the relation Eq. (6), we can write

$$\sum_{i=1}^{N} a_i v_i = \frac{\Delta}{2}\sum_{i=1}^{N} v_i^2 \tag{19}$$

Squaring both sides and substituting into Eq. (18), we arrive at the key identity

$$S'_C = \frac{1}{\sum_{i=1}^{N} v_i^2}\left(\sum_{i=1}^{N} v_i a_i\right)^2 = \frac{(\mathbf{v}\cdot\mathbf{a})^2}{\|\mathbf{v}\|^2} \tag{20}$$

This result expresses the transformed cosine similarity entirely in terms of the projection of **a** onto the direction defined by **v**.

3.6. Diagonalization and the Unitary Transformation

Since **H′** is real and symmetric, it can be diagonalized by an orthogonal matrix **U** whose columns are the eigenvectors of **H′**, with the first column equal to $\hat{\mathbf{v}}$

$$\mathbf{D} = \mathbf{U}^\dagger \mathbf{H}' \mathbf{U} = \mathrm{diag}(1, 0, \ldots, 0) \tag{21}$$

Defining the transformed vector $\mathbf{a}' = \mathbf{U}^\dagger\mathbf{a}$, and using the orthogonality $\mathbf{U}^T = \mathbf{U}^\dagger$ so that $\mathbf{U}\mathbf{D}\mathbf{U}^\dagger$ is symmetric, the transformed cosine similarity becomes

$$S'_C = \mathbf{a}^T\mathbf{H}'\mathbf{a} = \mathbf{a}^T(\mathbf{U}\mathbf{D}\mathbf{U}^\dagger)\mathbf{a} = (\mathbf{U}^\dagger\mathbf{a})^T\mathbf{D}(\mathbf{U}^\dagger\mathbf{a}) = \mathbf{a}'^T\mathbf{D}\mathbf{a}' \tag{22}$$

Writing this explicitly in matrix form

$$S'_C = \begin{bmatrix} a'_1 & a'_2 & \cdots & a'_N \end{bmatrix} \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} a'_1 \\ a'_2 \\ \vdots \\ a'_N \end{bmatrix} = a'^2_1 \tag{23}$$

where $a'_1 = \hat{\mathbf{v}}\cdot\mathbf{a}$ is the projection of **a** onto the bisector direction, consistent with Eq. (10). The unitary transformation preserves the norm

$$\sum_{i=1}^{N} a'^2_i = \|\mathbf{a}'\|^2 = \|\mathbf{a}\|^2 = 1 \tag{24}$$

where the equality $\|\mathbf{a}'\| = \|\mathbf{a}\|$ follows from the unitarity of **U**, and $\|\mathbf{a}\| = 1$ from L2 normalization. The remaining components $a'_2, \ldots, a'_N$ therefore encode the part of **a** orthogonal to $\hat{\mathbf{v}}$ and do not contribute to the similarity, but are required to completely specify the state of the system. The matrix **U** is constructed by taking $\hat{\mathbf{v}}$ as the first column and completing it to an orthonormal basis for $\mathbb{R}^N$, for example via QR decomposition; the choice of the remaining columns does not affect S'_C since the corresponding eigenvalues are zero.

The result $S'_C = a'^2_1$ is the central equation of the classical system. It expresses the transformed cosine similarity as the square of a single projection in the eigenbasis of the rank-1 Hamiltonian **H′**, and it is this form that admits a direct and exact quantum mechanical interpretation, as we show in Section IV.

3.7. The Degenerate Case a = −b

When **a** = −**b**, the cosine similarity is $S_C = -1$ and $S'_C = 0$. In this case **a** + **b** = 0 and the bisector $\hat{\mathbf{v}}$ is undefined. The result $S'_C = 0$ is nevertheless consistent with the framework: since $N \geq 2$, there exist unit vectors orthogonal to **a**, and any such $\hat{\mathbf{v}}$ satisfies $(\hat{\mathbf{v}}\cdot\mathbf{a})^2 = 0$ and $(\hat{\mathbf{v}}\cdot\mathbf{b})^2 = (\hat{\mathbf{v}}\cdot(-\mathbf{a}))^2 = 0$, so the Hamiltonian $\mathbf{H}' = \hat{\mathbf{v}}\hat{\mathbf{v}}^T$ yields $\mathbf{a}^T\mathbf{H}'\mathbf{a} = 0 = S'_C$ regardless of which orthogonal direction is chosen.

The degenerate case therefore does not affect the validity of the framework; it simply requires a different choice of $\hat{\mathbf{v}}$, and the result $S'_C = 0$ is recovered correctly for any such choice.

4. The LLM Embedding Quantum System

We now show that the classical system admits an exact quantum mechanical representation. The transition from the classical to the quantum formulation requires two conceptual steps: the promotion of real vector components to complex, time-dependent amplitudes, and the promotion of the Hamiltonian matrix to an operator acting on a Hilbert space. We show that these steps preserve all physically observable quantities—in particular, the transformed cosine similarity S'_C—and that the quantum formulation is not an approximation but a structurally exact equivalent of the classical system.

4.1. The Classical-to-Quantum Mapping and State Construction

There are two fundamental characteristics that distinguish a quantum mechanical system from a classical one: the state vector is complex-valued, which introduces the concept of phase, and the state vector is time-dependent, with its evolution governed by the Schrödinger equation. We now construct the quantum analogue of the classical embedding system by introducing both of these features through a canonical mapping.

4.1.1. The Canonical Mapping

In the classical system, the embedding vector $\mathbf{a} = [a_1, a_2,...,a_N]^T$ is a real, static vector, and the Hamiltonian $\mathbf{H'}$ is a real, symmetric matrix. To construct the quantum system, we promote these objects as follows:

1. The real vector $\mathbf{a}$ is promoted to a complex, time-dependent state vector $|\Psi(t)\rangle$ in an N-dimensional Hilbert space $\mathcal{H}$. Its initial condition is $|\Psi(0)\rangle = \mathbf{a}$ (where $\mathbf{a}$ is interpreted as a column vector in the computational basis).
2. The real symmetric matrix $\mathbf{H'}$ is promoted to a Hermitian operator $\hat{\mathbf{H}}'$ acting on $\mathcal{H}$. In the computational basis, $\hat{\mathbf{H}}'$ has the same matrix representation as $\mathbf{H'}$.
3. Similarly, the diagonalized classical vector $\mathbf{a'}$ is promoted to a complex, time-dependent state vector $|\tilde{\psi}(t)\rangle$ in the eigenbasis of $\hat{\mathbf{H}}'$. Its initial condition is $|\tilde{\psi}(0)\rangle = \mathbf{a}'$.

Since $\mathbf{H'}$ is already real and symmetric, it is automatically Hermitian, and the operator $\hat{\mathbf{H}}'$ has the same matrix representation as $\mathbf{H'}$ in the computational basis. The distinction between $\mathbf{H'}$ and $\hat{\mathbf{H}}'$ is therefore one of interpretation: in the classical system, $\mathbf{H'}$ acts on real vectors by matrix multiplication; in the quantum system, $\hat{\mathbf{H}}'$ acts on state vectors in a Hilbert space and its eigenvalues correspond to observable quantities.

4.1.2. Construction of the Quantum State in the Eigenbasis

Let $\{|n\rangle\}_{n=1}^{N}$ denote the eigenvectors of $\hat{\mathbf{H}}'$, which form an orthonormal basis for $\mathcal{H}$: $\langle m \,|\, n\rangle = \delta_{mn}$. The eigenvalue equation for $\hat{\mathbf{H}}'$ is

$$\hat{\mathbf{H}}'|n\rangle = E_n|n\rangle \tag{25}$$

where, from the rank-1 structure, the eigenvalues are $E_1 = 1$ and $E_n = 0$ for $n = 2,3,..., N$. The quantum state vector in the eigenbasis, $|\tilde{\psi}(t)\rangle$, is expressed as a superposition of these eigenvectors. We define its time-dependent coefficients as

$$\tilde{\psi}_n(t) = \tilde{A}_n \, e^{-iE_n t/\hbar} \tag{26}$$

Here $\tilde{A}_n$ is a real amplitude, E_n is the n-th eigenvalue of $\hat{\mathbf{H}}'$, t is the time parameter, $\hbar$ is a scaling constant introduced for dimensional consistency, and i is the imaginary unit. The real amplitudes $\tilde{A}_n$ are determined by the initial condition: at $t = 0$, the quantum state must reproduce the classical embedding vector in the eigenbasis. Specifically, we require

$$\tilde{\psi}_n(0) = \tilde{A}_n = a'_n \tag{27}$$

where a'_n are the components of the transformed classical vector $\mathbf{a'} = \mathbf{U}^{\dagger}\mathbf{a}$. This identification is natural: the eigenvectors $|n\rangle$ of $\hat{\mathbf{H}}'$ correspond to the columns of the unitary matrix $\mathbf{U}$, and the components a'_n are precisely the projections of a onto these eigenvectors. Thus, the quantum state in the eigenbasis is

$$|\tilde{\psi}(t)\rangle = \sum_{n=1}^{N} \tilde{\psi}_n(t)|n\rangle \tag{28}$$

The state in the original computational basis, $|\Psi(t)\rangle$, is related to $|\tilde{\psi}(t)\rangle$ by the unitary transformation $\mathbf{U}$: $|\Psi(t)\rangle = \mathbf{U}|\tilde{\psi}(t)\rangle$.

4.1.3. The Role of the Complex Phase and Normalization

The key difference between the classical vector $\mathbf{a'}$ and the quantum state $|\tilde{\psi}(t)\rangle$ is the presence of the complex phase factors $e^{-iE_n t/\hbar}$. At $t = 0$, these phase factors are all equal to unity, and the quantum state in the eigenbasis reduces to the classical vector $\mathbf{a'}$

$$|\tilde{\psi}(0)\rangle = \sum_{n=1}^{N} a'_n |n\rangle \quad \longleftrightarrow \quad \mathbf{a}' = [a'_1, a'_2, \ldots, a'_N]^T \tag{29}$$

For $t \neq 0$, the individual coefficients $\tilde{\psi}_n(t)$ acquire different phases depending on their eigenvalues E_n. However, the squared moduli of the coefficients are time-independent

$$|\tilde{\psi}_n(t)|^2 = |\tilde{A}_n e^{-iE_n t/\hbar}|^2 = \tilde{A}_n^2 = a_n'^2 \tag{30}$$

This is a crucial result. The complex phase factors, which encode the time evolution of the quantum state, do not affect the squared amplitudes. Since all observable quantities in the framework—in particular, the transformed cosine similarity S'_C —depend only on squared amplitudes, the quantum system yields identical physical predictions to the classical system at all times, not only at $t = 0$.

The normalization of the quantum state is preserved by construction. At any time t

$$\langle \tilde{\psi}(t) \,|\, \tilde{\psi}(t) \rangle = \sum_{n=1}^{N} |\tilde{\psi}_n(t)|^2 = \sum_{n=1}^{N} \tilde{A}_n^2 = \sum_{n=1}^{N} a_n'^2 = \|\mathbf{a}'\|^2 = \|\mathbf{a}\|^2 = 1 \tag{31}$$

where the last two equalities follow from the unitarity of $\mathbf{U}$ and the L2 normalization of $\mathbf{a}$. The quantum state $|\tilde{\psi}(t)\rangle$ therefore lives on the unit sphere in $\mathcal{H}$ at all times, consistent with the Born rule interpretation of $|\tilde{\psi}_n(t)|^2$ as probabilities.

4.2. The Time-Dependent Schrödinger Equation and Its Solutions

We now demonstrate that the time evolution defined by the coefficients $\tilde{\psi}_n(t)$ in Eq. (26) satisfies the time-dependent Schrödinger equation in both the eigenbasis and the computational basis. This establishes that the quantum embedding system obeys the fundamental dynamical law of quantum mechanics.

4.2.1. Verification of the Schrödinger Equation in the Eigenbasis

First, we verify that the time evolution of $|\tilde{\psi}(t)\rangle$ (as defined in Eq. (28) with coefficients from Eq. (26)) satisfies the Schrödinger equation in the eigenbasis. Let $\hat{\mathbf{D}}$ be the diagonal operator representing the Hamiltonian in this basis, with $\hat{\mathbf{D}}|n\rangle = E_n|n\rangle$.

We take the time derivative of $|\tilde{\psi}(t)\rangle$

$$i\hbar \frac{d}{dt}|\tilde{\psi}(t)\rangle = i\hbar \sum_{n=1}^{N} \frac{d\tilde{\psi}_n(t)}{dt}|n\rangle = i\hbar \sum_{n=1}^{N} \tilde{A}_n \left(-\frac{iE_n}{\hbar}\right) e^{-iE_n t/\hbar}|n\rangle = \sum_{n=1}^{N} E_n \tilde{\psi}_n(t)|n\rangle \tag{32}$$

On the other hand, applying the diagonal Hamiltonian operator $\hat{\mathbf{D}}$ to $|\tilde{\psi}(t)\rangle$

$$\hat{\mathbf{D}}|\tilde{\psi}(t)\rangle = \sum_{n=1}^{N} \tilde{\psi}_n(t)\hat{\mathbf{D}}|n\rangle = \sum_{n=1}^{N} E_n \tilde{\psi}_n(t)|n\rangle \tag{33}$$

Since Eqs. (32) and (33) are identical, the quantum state in the eigenbasis satisfies the time-dependent Schrödinger equation

$$i\hbar \frac{d}{dt}|\tilde{\psi}(t)\rangle = \hat{\mathbf{D}}|\tilde{\psi}(t)\rangle \tag{34}$$

This is the quantum dynamical equation governing the LLM embedding quantum system in its eigenbasis. The diagonal Hamiltonian $\hat{\mathbf{D}}$ plays the role of the energy operator, and its eigenvalues {1,0,...,0} determine the time evolution of each component of the state vector.

4.2.2. Consistency with the Schrödinger Equation in the Computational Basis

Now we show that if $|\tilde{\psi}(t)\rangle$ satisfies Eq. (34), then the state in the original computational basis, $|\Psi(t)\rangle = \mathbf{U}|\tilde{\psi}(t)\rangle$, satisfies its corresponding Schrödinger equation with $\hat{\mathbf{H}}'$. Starting from Eq. (34)

$$i\hbar\frac{d}{dt}|\tilde{\psi}(t)\rangle = \hat{\mathbf{D}}|\tilde{\psi}(t)\rangle \tag{35}$$

Multiply both sides by $\mathbf{U}$ from the left

$$i\hbar\,\mathbf{U}\frac{d}{dt}|\tilde{\psi}(t)\rangle = \mathbf{U}\hat{\mathbf{D}}|\tilde{\psi}(t)\rangle \tag{36}$$

Since $\mathbf{U}$ is time-independent, we can move it inside the derivative on the left side

$$i\hbar\frac{d}{dt}\left(\mathbf{U}|\tilde{\psi}(t)\rangle\right) = \mathbf{U}\hat{\mathbf{D}}|\tilde{\psi}(t)\rangle \tag{37}$$

Substitute $|\Psi(t)\rangle = \mathbf{U}|\tilde{\psi}(t)\rangle$ on the left side. On the right side, use the diagonalization relation $\hat{\mathbf{D}} = \mathbf{U}^{\dagger}\hat{\mathbf{H}}'\mathbf{U}$, which implies $\mathbf{U}\hat{\mathbf{D}} = \hat{\mathbf{H}}'\mathbf{U}$

$$i\hbar\frac{d}{dt}|\Psi(t)\rangle = \hat{\mathbf{H}}'\left(\mathbf{U}|\tilde{\psi}(t)\rangle\right) \tag{38}$$

Then substitute $|\Psi(t)\rangle = \mathbf{U}|\tilde{\psi}(t)\rangle$ on the right side

$$i\hbar\frac{d}{dt}|\Psi(t)\rangle = \hat{\mathbf{H}}'|\Psi(t)\rangle \tag{39}$$

This confirms that the state $|\Psi(t)\rangle$ in the computational basis also satisfies the Schrödinger equation with the original Hamiltonian operator $\hat{\mathbf{H}}'$.

4.2.3. Matrix Representations

To provide clarity to the notations and results, we also show the matrix forms. In matrix representation, with $\Psi_j(t)$ as the components of $|\Psi(t)\rangle$ in the computational basis, the Eq. (39) reads

$$i\hbar\frac{d}{dt}\begin{bmatrix}\Psi_1(t)\\ \Psi_2(t)\\ \vdots\\ \Psi_N(t)\end{bmatrix} = \frac{1}{\|\mathbf{v}\|^2}\begin{bmatrix}v_1^2 & v_1v_2 & \cdots & v_1v_N\\ v_2v_1 & v_2^2 & \cdots & v_2v_N\\ \vdots & \vdots & \ddots & \vdots\\ v_Nv_1 & v_Nv_2 & \cdots & v_N^2\end{bmatrix}\begin{bmatrix}\Psi_1(t)\\ \Psi_2(t)\\ \vdots\\ \Psi_N(t)\end{bmatrix} \tag{40}$$

Similarly, the time-dependent Schrödinger equation in the eigenbasis is

$$i\hbar\frac{d}{dt}\begin{bmatrix}\tilde{\psi}_1(t)\\ \tilde{\psi}_2(t)\\ \vdots\\ \tilde{\psi}_N(t)\end{bmatrix} = \begin{bmatrix}1 & 0 & \cdots & 0\\ 0 & 0 & \cdots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \cdots & 0\end{bmatrix}\begin{bmatrix}\tilde{\psi}_1(t)\\ \tilde{\psi}_2(t)\\ \vdots\\ \tilde{\psi}_N(t)\end{bmatrix} \tag{41}$$

4.2.4. Equivalence with the Classical System and Observables

The quantum mechanical expectation value of the Hamiltonian operator provides the link to the classical transformed cosine similarity. The expectation value of $\hat{\mathbf{D}}$ in the state $|\tilde{\psi}(t)\rangle$ is

$$\langle\tilde{\psi}(t)|\hat{\mathbf{D}}|\tilde{\psi}(t)\rangle = \sum_{n=1}^{N} E_n|\tilde{\psi}_n(t)|^2 = E_1|\tilde{\psi}_1(t)|^2 = |\tilde{\psi}_1(t)|^2 = \tilde{A}_1^2 = a_1'^2 = S_C' \tag{42}$$

where we have used $E_1 = 1$, $E_n = 0$ for $n \geq 2$, and the time-independence of the squared amplitudes established in Eq. (30). Similarly, the expectation value of $\hat{\mathbf{H}}'$ in the state $|\Psi(t)\rangle$ is

$$\langle\Psi(t)|\hat{\mathbf{H}}'|\Psi(t)\rangle = \langle\mathbf{U}\tilde{\psi}(t)|\mathbf{U}\hat{\mathbf{D}}\mathbf{U}^\dagger|\mathbf{U}\tilde{\psi}(t)\rangle = \langle\tilde{\psi}(t)|\hat{\mathbf{D}}|\tilde{\psi}(t)\rangle = S'_C \tag{43}$$

Both expectation values are identical to the classical transformed cosine similarity at all times. This confirms that the Schrödinger equation, despite introducing complex phases and time dependence, preserves the physical content of the classical system.

4.2.5. Geometric Interpretation

The relationship between the real amplitude a'_n and the complex coefficient $\tilde{\psi}_n(t)$ has a clear geometric interpretation. Consider the complex plane for each component $\tilde{\psi}_n(t)$. At $t = 0$, the coefficient $\tilde{\psi}_n(0) = a'_n$ lies on the real axis at distance $|a'_n|$ from the origin. As time evolves:

• For $n = 1$ (the excited state with $E_1 = 1$): the coefficient $\tilde{\psi}_1(t) = \tilde{A}'_1\, e^{-it/\hbar}$ traces a circle of radius $|a'_1|$ in the complex plane, with angular velocity $-1/\hbar$. The squared modulus $|\tilde{\psi}_1(t)|^2 = a'^2_1$ is the squared radius of this circle, which is constant.
• For $n \geq 2$ (the ground states with $E_n = 0$): the coefficients $\tilde{\psi}_n(t) = \tilde{A}'_n$ remain fixed on the real axis. There is no phase evolution because the eigenvalue is zero.

5. The Commutative Diagram and Structural Isomorphism

We now collect the results of Sections III and IV into a unified picture. The transformed cosine similarity S'_C admits four equivalent expressions, organized by two independent operations—diagonalization and quantization:

$$\text{Classical, original basis:} \qquad S'_C = \mathbf{a}^T\mathbf{H}'\mathbf{a} \tag{44}$$

$$\text{Classical, eigenbasis:} \qquad S'_C = \mathbf{a}'^T\mathbf{D}\mathbf{a}' = a'^2_1 \tag{45}$$

$$\text{Quantum, original basis:} \qquad S'_C = \langle\Psi(t)|\hat{\mathbf{H}}'|\Psi(t)\rangle \tag{46}$$

$$\text{Quantum, eigenbasis:} \qquad S'_C = \langle\tilde{\psi}(t)|\hat{\mathbf{D}}|\tilde{\psi}(t)\rangle = |\tilde{\psi}_1|^2 \tag{47}$$

All four expressions yield the same numerical value for any pair of embedding vectors **a** and **b**. The transformations connecting them are shown in Figure 1.

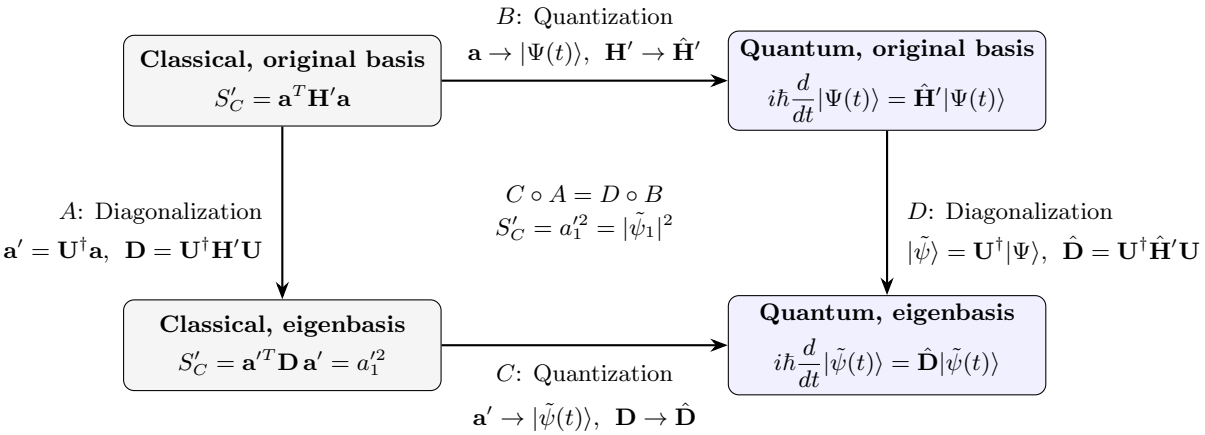

Figure 1: Commutative Diagram Relating the Four Equivalent Formulations of the Transformed Cosine Similarity S'_C . The boxes on left denote classical formulations; the boxes on right denote quantum formulations. Arrow A: classical diagonalization via unitary transformation **U**. Arrow B: quantization in the original basis, promoting real vectors to complex state vectors and symmetric matrices to Hermitian operators. Arrow C: quantization in the eigenbasis. Arrow D: diagonalization of the Schrödinger equation by the same unitary **U**. The diagram commutes, $C \circ A = D \circ B$, meaning that diagonalization and quantization can be performed in either order with identical results.

Each arrow corresponds to a well-defined mathematical operation. Arrow A is the classical change of basis $\mathbf{a}' = \mathbf{U}^\dagger\mathbf{a}$, $\mathbf{D} = \mathbf{U}^\dagger\mathbf{H'U}$. Arrows B and C promote real vectors to complex state vectors with coefficients $\tilde{\psi}_n(t) = \tilde{A}'_n e^{-iE_n t/\hbar}$ and real symmetric matrices to Hermitian operators, so that bilinear forms become expectation values. Arrow D is the quantum change of basis $|\tilde{\psi}(t)\rangle = \mathbf{U}^\dagger|\Psi(t)\rangle$. The commutativity $C \circ A = D \circ B$ follows from the time-independence of **U**, which ensures that the unitary transformation commutes with the time derivative and therefore with the Schrödinger equation. The result is independent of the order in which diagonalization and quantization are performed, and is summarized by the equivalence chain

$$S'_C = \mathbf{a}^T \mathbf{H}' \mathbf{a} = a'^2_1 = |\tilde{\psi}_1(t)|^2 = \langle \Psi(t)|\hat{\mathbf{H}}'|\Psi(t)\rangle \tag{48}$$

in which the first two equalities are classical and the last two are quantum mechanical.

5.1. Nature of the Isomorphism

The commutative diagram establishes a structural isomorphism between the classical and quantum formulations in the following precise sense. Algebraic structure is preserved: vector space operations, inner products, and operator actions map consistently between $\mathbb{R}^N$ and $\mathcal{H}$. Observable quantities are identical: S'_C takes the same value in all four representations. The mapping is invertible: given $|\Psi(t)\rangle$ at any time, the classical vector a can be recovered from the squared amplitudes $|\tilde{\psi}_n(t)|^2 = a'^2_n$ and the known eigenvectors of $\hat{\mathbf{H}}'$. Transformation laws are consistent: orthogonal transformations in $\mathbb{R}^N$ map to unitary transformations in $\mathcal{H}$, both preserving norms and commuting with diagonalization.

The isomorphism is *structural*, not *physical*. The mathematical structures used to describe LLM embeddings—real vector spaces, symmetric bilinear forms, orthogonal transformations, and squared projections—are isomorphic to those of quantum mechanics—Hilbert spaces, Hermitian operators, unitary transformations, and Born rule probabilities. The empirically verifiable content of this isomorphism is the identity

$$S'_C = a'^2_1 = |\tilde{\psi}_1(t)|^2 = P(|0\ldots0\rangle) \tag{49}$$

where the first term is a classical squared projection, the second is a quantum squared amplitude, and the third is a measurement probability on a quantum computer. The left-hand side is computed classically from LLM embeddings; the right-hand side is measured on quantum hardware. Their equality is confirmed experimentally in Section VIII.

5.2. Additional Structure in the Quantum Formulation

While the classical and quantum systems yield identical values of S'_C, the quantum formulation provides structure with no classical counterpart. The squared amplitudes $\{|\tilde{\psi}_n|^2\}$ form a probability distribution over eigenstates of $\hat{\mathbf{H}}'$, decomposing the embedding into a similarity-aligned component $|\tilde{\psi}_1|^2 = a'^2_1$ and orthogonal components $|\tilde{\psi}_2|^2, \ldots, |\tilde{\psi}_N|^2$ that constitute the semantic noise profile of the embedding pair [14]. The complex phases $e^{-iE_n t/\hbar}$ are invisible in the diagonal observable S'_C but become observable in off-diagonal measurements and in the gauge field extension of the framework [18]. Most significantly, the quantum formulation admits a local U(1) gauge symmetry under $\tilde{\psi}_1(t) \to e^{i\theta(t)}\tilde{\psi}_1(t)$, with an associated conserved semantic charge and a gauge field $A_0(t)$ modeling contextual influence on the embedding [18]. This symmetry structure has no counterpart in the classical description and is accessible only through the quantum formulation.

The classical system is embedded within the quantum system as the special case where all phases are zero and measurements are performed in the eigenbasis. The quantum formulation is therefore strictly richer than the classical one, while remaining fully consistent with it at the level of all diagonal observables.

6. Interpretation of Time and the Complex State Vector

The quantum formulation introduces two features absent from the classical description: a complex state vector with phase information, and a time parameter t governing its evolution. Since the embedding vectors **a** and **b** are static objects produced by the LLM encoder, the natural question is: what does t mean in a context where there is no obvious physical time?

The resolution is that t is a gauge parameter. In the relation $\hat{U}(t) = e^{-i\hat{H}t/\hbar}$, only the product $\hat{H}t/\hbar$ is physically meaningful; rescaling $t \to \alpha t$ and $\hat{H} \to \hat{H}/\alpha$ leaves $\hat{U}$ unchanged. The unitary transformation **U** is fixed by the embedding vectors through the eigenvectors of $\hat{\mathbf{H}}'$, and this fixes the gauge-invariant product $\hat{\mathbf{H}}'t/\hbar$ without determining t independently. The time parameter is therefore not an observable quantity in the LLM context; it is an internal degree of freedom of the unitary evolution that cancels out of all observable predictions. We make this precise in the following subsections, using the analogy with quantum circuit gates as a concrete illustration.

6.1. The Time Evolution Operator and Its Role

The formal solution of the Schrödinger equation for the state in the original basis, $i\hbar\frac{d}{dt}|\Psi(t)\rangle = \hat{\mathbf{H}}'|\Psi(t)\rangle$, is

$$|\Psi(t)\rangle = e^{-i\hat{\mathbf{H}}'t/\hbar}|\Psi(0)\rangle \equiv \hat{U}(t)|\Psi(0)\rangle \tag{50}$$

where $\hat{U}(t) = e^{-i\hat{\mathbf{H}}'t/\hbar}$ is the time evolution operator, unitary for all real t. The time parameter enters the formalism solely through this operator, parameterizing a one-parameter family of unitary transformations generated by $\hat{\mathbf{H}}'$. When expressed in the eigenbasis, the state is $|\tilde{\psi}(t)\rangle = \mathbf{U}^\dagger|\Psi(t)\rangle$, and its evolution is given by $|\tilde{\psi}(t)\rangle = \hat{U}(t)_{\text{diag}}|\tilde{\psi}(0)\rangle$, where $\hat{U}(t)_{\text{diag}} = \mathbf{U}^\dagger\hat{U}(t)\mathbf{U}$. For the rank-1 Hamiltonian of the embedding system, the diagonalized evolution operator takes a particularly simple form in the eigenbasis

$$\hat{U}(t)_{\text{diag}} = \text{diag}(e^{-it/\hbar}, 1, 1, \ldots, 1) \tag{51}$$

Only the first component acquires a phase; all others are unchanged. The entire effect of time evolution is a single phase rotation in the one-dimensional subspace spanned by $\hat{\mathbf{v}}$. The squared modulus of every component of $|\tilde{\psi}(t)\rangle$ is therefore time-independent

$$|\tilde{\psi}_n(t)|^2 = |\tilde{A}'_n e^{-iE_n t/\hbar}|^2 = a'^2_n \quad \text{for all } t \tag{52}$$

and in particular $S'_C = |\tilde{\psi}_1(t)|^2 = a'^2_1$ at all times. The time parameter is present in the formalism but absent from every observable prediction.

6.2. The Gauge Parameter Interpretation

The non-uniqueness of the decomposition $\hat{U} = e^{-i\hat{H}t/\hbar}$ is the precise sense in which t is a gauge parameter. Given a fixed unitary $\hat{U}$, there is a one-parameter family of decompositions indexed by $t > 0$

$$\hat{U} = e^{-i\hat{H}(t)\cdot t/\hbar}, \qquad \hat{H}(t) = \frac{\hbar}{t}\hat{H}_0 \tag{53}$$

where $\hat{H}_0$ is any fixed generator of $\hat{U}$. Different values of t give different Hamiltonians but the same unitary and therefore the same observable predictions. The physically meaningful object is the generator $\hat{H}_0 = \hat{H}(t) \cdot t/\hbar$, not t or $\hat{H}(t)$ separately. In the LLM embedding system, $\hat{H}_0$ is fixed by the embedding vectors through $\hat{\mathbf{H}}' = \hat{\mathbf{v}}\hat{\mathbf{v}}^T$, and the diagonalizing unitary U is determined by the eigenvectors of $\hat{\mathbf{H}}'$. The gauge parameter t is the remaining freedom, and it drops out of all observables. This is the complete resolution of the time interpretation question: asking what t means physically is analogous to asking what the absolute value of the electromagnetic potential means physically—it is the wrong question, because its specific value can be arbitrarily chosen without affecting any observable outcome.

6.3. The Hadamard Gate as Schrödinger Time Evolution

To make the gauge parameter interpretation concrete, we work through a complete example using the Hadamard gate which is a well-known gate in quantum computing. We find a Hamiltonian $\mathcal{H}$ and a time t such that $e^{-i\mathcal{H}t/\hbar}$ exactly reproduces the Hadamard gate, demonstrating explicitly that the time parameter is absorbed into the gate definition and does not appear as an independent observable. This example also illustrates the rank-1 Hamiltonian structure shared with the LLM embedding system.

6.3.1. The Hadamard Gate and its Eigenstructure

The Hadamard gate is defined by

$$H = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{54}$$

Its eigenvalues follow from the characteristic equation $\det(H - \lambda\mathbf{I}) = 0$

$$\lambda^2 = 1 \quad \Rightarrow \quad \lambda_1 = +1, \quad \lambda_2 = -1 \tag{55}$$

The corresponding normalized eigenvectors are

$$|v_1\rangle = \frac{1}{\sqrt{2+\sqrt{2}}}\begin{bmatrix}1+\frac{1}{\sqrt{2}}\\ \frac{1}{\sqrt{2}}\end{bmatrix} \quad (\lambda_1 = +1) \tag{56}$$

$$|v_2\rangle = \frac{1}{\sqrt{2-\sqrt{2}}}\begin{bmatrix}-1+\frac{1}{\sqrt{2}}\\ \frac{1}{\sqrt{2}}\end{bmatrix} \quad (\lambda_2 = -1) \tag{57}$$

These eigenvectors are orthonormal, $\langle v_1 \mid v_2\rangle = 0$, and the Hadamard gate admits the spectral decomposition

$$H = (+1)|v_1\rangle\langle v_1| + (-1)|v_2\rangle\langle v_2| = |v_1\rangle\langle v_1| - |v_2\rangle\langle v_2| \tag{58}$$

6.3.2. Constructing the Hamiltonian

We seek $\mathcal{H}$ and t such that $H = e^{-i\mathcal{H}t/\hbar}$. Writing the time evolution operator in the eigenbasis of $\mathcal{H}$

$$e^{-i\mathcal{H}t/\hbar} = \sum_k e^{-i\epsilon_k t/\hbar}|w_k\rangle\langle w_k| \tag{59}$$

and comparing with Eq. (58), we identify $|w_k\rangle = |v_k\rangle$ and require

$$e^{-i\epsilon_1 t/\hbar} = +1 \tag{60}$$

$$e^{-i\epsilon_2 t/\hbar} = -1 \tag{61}$$

These are satisfied by $\epsilon_1 t/\hbar = 2\pi n$ and $\epsilon_2 t/\hbar = (2m+1)\pi$ for integers *n,m*. The simplest solution (*n* = 0, *m* = 0) gives

$$\epsilon_1 = 0, \qquad \epsilon_2 = \frac{\pi\hbar}{t} \tag{62}$$

and the Hamiltonian

$$\mathcal{H} = \frac{\pi\hbar}{t}|v_2\rangle\langle v_2| \tag{63}$$

This has a rank-1 structure: only $|v_2\rangle$ has nonzero energy, and $|v_1\rangle$ does not evolve in time. The gauge parameter nature of *t* is explicit: $\mathcal{H}$ depends on *t* through the prefactor $\pi\hbar/t$, so any $t > 0$ yields a valid decomposition with a correspondingly rescaled Hamiltonian. The physically meaningful object is the product $\mathcal{H} \cdot t/\hbar = \pi|v_2\rangle\langle v_2|$, which is independent of *t*.

6.3.3. Verification

Direct computation confirms $e^{-i\mathcal{H}t/\hbar} = H$

$$e^{-i\mathcal{H}t/\hbar} = e^0|v_1\rangle\langle v_1| + e^{-i\pi}|v_2\rangle\langle v_2| = |v_1\rangle\langle v_1| - |v_2\rangle\langle v_2| = H \tag{64}$$

When the Hadamard gate appears in a circuit diagram, only the symbol *H* is drawn. The time *t* and Hamiltonian $\mathcal{H}$ are internal to the gate implementation; a user of the circuit neither knows nor needs their values. The measurement output is a real probability; the complex phases generated during time evolution are not directly observed but contribute to the measurement probability through the Born rule.

6.3.4. Connection to the LLM Embedding System

The Hadamard example illustrates the general principle governing the LLM embedding quantum system. In both cases:

- A unitary transformation is generated by a rank-1 Hamiltonian acting for some time *t*.
- The time parameter is a gauge variable: only the product $\mathcal{H}t/\hbar$ is physically meaningful, and this product is fixed by the unitary.
- The Hamiltonian has a single nonzero eigenvalue, so time evolution consists of a single phase rotation in a one-dimensional subspace.
- The observable output is a real-valued probability obtained by squaring the relevant amplitude, in which the phase cancels.

The difference is one of context, not mathematical structure. In a quantum circuit, the Hamiltonian and time are determined by the physical implementation of the gate. In the LLM embedding system, the Hamiltonian is determined by the embedding vectors through $\hat{\mathbf{H}}' = \hat{\mathbf{v}}\hat{\mathbf{v}}^T$, and *t* is the gauge freedom that does not correspond to any physical duration. The absence of an explicit time variable does not invalidate the quantum description; quantum circuits demonstrate that fully quantum mechanical systems can be described,

implemented, and measured without ever referencing time directly.

6.3.5. The Complex State Vector as a Theoretical Construct

In the LLM embedding quantum system, the complex state vector (whether $|\Psi(t)\rangle$ in the computational basis or $|\tilde{\psi}(t)\rangle$ in the eigenbasis) serves three essential purposes. First, it ensures unitary evolution: the Schrödinger equation requires complex amplitudes to generate the phase factors $e^{-iE_n t/\hbar}$ that make the time evolution operator unitary, since real-valued evolution equations do not in general preserve the norm under non-trivial Hamiltonians. Second, it encodes phase relationships: the relative phases between components of $|\tilde{\psi}(t)\rangle$ carry information about the Hamiltonian eigenstructure and become observable when the system is measured in a rotated basis, as occurs in quantum circuit implementations and in the gauge field extension. Third, it enables quantum circuit realization: quantum computers operate natively on complex state vectors, and the state preparation, unitary transformation, and measurement steps all rely on the complex Hilbert space structure. The complex state vector is therefore not an arbitrary mathematical decoration but the minimal structure required to simultaneously satisfy the Schrödinger equation, preserve unitarity, and enable quantum circuit implementation.

6.4. Interpretation of $\hbar$ as Semantic Resolution

The parameter $\hbar$ plays a role symmetric to t: both are individually unobservable, and only their ratio

$$\theta_s \equiv \frac{t}{\hbar} \tag{65}$$

enters the time evolution operator. We call θ_s the *semantic phase angle*. In terms of θ_s, the time evolution operator Eq. (51) takes the form

$$\hat{U}_{\text{diag}} = \text{diag}(e^{-i\theta_s}, 1, 1, \ldots, 1) \tag{66}$$

and the coefficients Eq. (26) become

$$\tilde{\psi}_1(\theta_s) = \tilde{A}'_1\, e^{-i\theta_s}, \qquad \tilde{\psi}_n(\theta_s) = \tilde{A}'_n \quad \text{for } n \geq 2 \tag{67}$$

The squared moduli $|\tilde{\psi}_n(\theta_s)|^2 = a'^2_n$ are independent of θ_s, confirming that the semantic phase angle, like t and $\hbar$ individually, does not enter any diagonal observable. The gauge freedom of the system is therefore two-dimensional: the simultaneous rescaling

$$t \;\to\; \alpha t, \qquad \hbar \;\to\; \alpha\hbar, \qquad \alpha > 0 \tag{68}$$

leaves θ_s, $\hat{U}_{\text{diag}}$, and all observable predictions invariant. The physically meaningful object is θ_s alone, not t or $\hbar$ separately. This is the precise sense in which $\hbar$ is a gauge parameter of the LLM embedding quantum system, symmetric in role to t.

The gauge freedom of the system is therefore two-dimensional. Section VI identified one redundancy: rescaling $t \to \alpha t$ and $\hat{\mathbf{H}}' \to \hat{\mathbf{H}}'/\alpha$ leaves $\hat{U}$ unchanged with $\hbar$ fixed. A second, independent redundancy is the simultaneous rescaling $t \to \alpha t$ and $\hbar \to \alpha\hbar$ with $\hat{\mathbf{H}}'$ fixed, which leaves $\theta_s = t/\hbar$ invariant. Together these confirm that neither t nor $\hbar$ is individually observable; the physically meaningful object is the product $\hat{\mathbf{H}}'t/\hbar$, of which $\theta_s = t/\hbar$ is the gauge-invariant part once $\hat{\mathbf{H}}'$ is fixed by the embedding vectors.

6.4.1. $\hbar$ as the Semantic Resolution Scale

Beyond its role as a gauge parameter, $\hbar$ admits a concrete interpretation in terms of the geometry of the embedding space. The embedding model maps a finite vocabulary of V tokens to unit vectors in $\mathbb{R}^N$. The minimum angular separation over all pairs of distinct token embeddings defines the *semantic resolution* of the model

$$\hbar_{\text{sem}} \equiv \min_{i \neq j} \arccos(\mathbf{a}_i \cdot \mathbf{a}_j) \tag{69}$$

This is the smallest angular distance at which the model can distinguish two distinct semantic units. Setting $\hbar = \hbar_{\text{sem}}$ provides a canonical gauge choice, making $\theta_s = t/\hbar_{\text{sem}}$ a dimensionless count of minimum semantic steps. This interpretation has three immediate consequences. First, $\hbar_{\text{sem}}$ is model-dependent: different LLMs have different semantic resolutions, reflecting their different capacities for semantic discrimination. Second, $\hbar_{\text{sem}}$ is consistent with the parameter Δ of Section III: both quantify the granularity of token representations, Δ in component space and $\hbar_{\text{sem}}$ on the unit sphere. Third, the interpretation yields a testable prediction: $\hbar_{\text{sem}}$ should decrease monotonically with embedding dimension N, since higher-dimensional models can represent finer semantic distinctions. This is measurable by computing $\min_{i \neq j} \arccos(\mathbf{a}_i \cdot \mathbf{a}_j)$ across a fixed vocabulary for models of varying dimension.

With $\hbar = \hbar_{sem}$ as the canonical gauge choice, the gauge field $A(\theta_s)$ of the U(1) symmetry [18] becomes dimensionless, measuring contextual perturbations in units of the minimum semantic step. As with the time parameter t, this interpretation is *structural* : the mathematical role of $\hbar$ in the formalism is identical to that of $\hbar_{sem}$ in the geometry of the embedding space, and the correspondence is exact at the level of mathematical structure.

6.4.2. Connection to the U (1) Gauge Structure

The semantic resolution scale connects naturally to the local U(1) gauge symmetry discussed in Section V and developed in Ref. [18]. Under the local gauge transformation

$$\tilde{\psi}_1(\theta_s) \rightarrow e^{i\alpha(\theta_s)}\tilde{\psi}_1(\theta_s) \tag{70}$$

the covariant derivative takes the form

$$D_{\theta_s} = \partial_{\theta_s} + iA(\theta_s) \tag{71}$$

where $A(\theta_s)$ is the gauge field modeling contextual influence on the embedding. With $\hbar = \hbar_{sem}$ as the canonical gauge choice, the gauge field $A(\theta_s)$ is dimensionless and measures contextual perturbations in units of the minimum semantic step. This gives the gauge field a concrete operational meaning: a unit perturbation $A(\theta_s) = 1$ shifts the semantic phase by one minimum semantic step, the smallest contextual influence that is distinguishable within the resolution of the embedding model.

6.4.3. Structural Interpretation

As with the time parameter t, the interpretation of $\hbar$ as $\hbar_{sem}$ is structural, not physical. We do not claim that LLM embedding spaces possess a physical analogue of Planck's constant, nor that the discreteness of the token vocabulary arises from quantum mechanical quantization. The claim is precisely that the mathematical role of $\hbar$ in the formalism—as the scale setting the ratio $\theta_s = t/\hbar$ that governs phase evolution—is identical to the mathematical role of the minimum angular separation $\hbar_{sem}$ in the geometry of the embedding space. The correspondence is exact at the level of mathematical structure, and it provides a canonical gauge fixing that connects the abstract parameter $\hbar$ to a directly computable property of any LLM embedding model.

7. Quantum Circuit Implementation

The structural isomorphism established in Sections III–IV shows that the transformed cosine similarity can be expressed as a quantum mechanical measurement probability. This section details a quantum circuit that computes S'_C from LLM embedding vectors, realizing the path from *B* to *D* in Figure 1. We first initialize the quantum circuit with the values of the classical embedding vector and then use a quantum computer to perform the diagonalization. As a result, we should be able to observe the same classical result S'_C .

The circuit operates on $n = \lceil \log_2 N \rceil$ qubits, where N is the embedding dimension (padded to the nearest power of 2), and consists of three stages: state preparation, unitary transformation, and measurement.

7.1. State Preparation

The first step encodes the classical L2-normalized embedding vector $\mathbf{a} = [a_1, a_2, ..., a_N]^T$ into an n-qubit quantum state

$$|\Psi(0)\rangle = \sum_{i=1}^{N} a_i |i\rangle \tag{72}$$

where $|i\rangle$ denotes the i-th computational basis state. The L2 normalization $\|\mathbf{a}\| = 1$ ensures $|\Psi(0)\rangle$ is a valid quantum state. If N is not a power of 2, a is zero-padded to 2^n dimensions, a process that preserves cosine similarity. This stage leverages quantum superposition to represent the N components of a using only n qubits.

7.2. Unitary Transformation

The second stage applies a unitary transformation $\mathbf{U}^\dagger$ to rotate the quantum state from the computational basis into the eigenbasis of the Hamiltonian $\hat{\mathbf{H}}'$. The matrix $\mathbf{U}$ is constructed such that its first column is the normalized angle bisector $\hat{\mathbf{v}} = \mathbf{v}/\|\mathbf{v}\|$, and the remaining columns form an orthonormal completion (e.g., via QR decomposition). This construction ensures that $\mathbf{U}^\dagger \hat{\mathbf{v}} = |0 \ldots 0\rangle$.

When $\mathbf{U}^\dagger$ acts on the prepared state $|\Psi(0)\rangle$, the state becomes $\mathbf{U}^\dagger|\Psi(0)\rangle = \sum_i a'_i|i\rangle$, where $a'_i = (\mathbf{U}^\dagger \mathbf{a})_i$. Crucially, the amplitude of the $|0 \ldots 0\rangle$ state in this transformed basis is $a'_1 = \hat{\mathbf{v}} \cdot \mathbf{a}$.

7.3. Measurement and Recovery of Cosine Similarity

The final stage measures all qubits in the computational basis. According to the Born rule, the probability of measuring a specific basis state $|i\rangle$ is $|a'_i|^2$. The outcome of primary interest is the all-zeros state $|0...0\rangle$

$$P(|0\ldots0\rangle) = |a'_1|^2 = (\hat{\mathbf{v}} \cdot \mathbf{a})^2 = S'_C \tag{73}$$

This is the central result: the probability of measuring $|0\ldots0\rangle$ directly yields the transformed cosine similarity S'_C. On a quantum computer, this probability is estimated by repeating the circuit N_{shots} times and calculating the frequency of the $|0\ldots0\rangle$ outcome. The original cosine similarity S_C can then be recovered using $S_C = 2P(|0\ldots0\rangle) - 1$.

7.4. The Complete Circuit

The full quantum circuit combines these three stages:

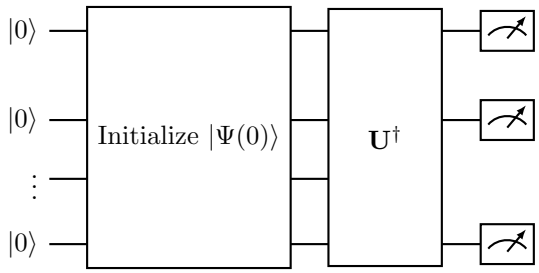

(74)

This circuit provides an experimentally verifiable realization of the structural isomorphism: a classical quantity (cosine similarity) is computed as a quantum measurement probability (Born rule). Unlike the general-purpose swap test, this circuit is specifically tailored to the LLM embedding similarity problem, directly yielding S'_C with fewer qubits and no ancilla. The framework's internal consistency is confirmed by the direct derivation: $\langle 0\ldots0|\mathbf{U}^\dagger|\Psi(0)\rangle = \hat{\mathbf{v}} \cdot \mathbf{a} = a'_1$, leading to $P(|0\ldots0\rangle) = a'^2_1 = S'_C$.

8. Experimental Validation

The theoretical framework predicts that the probability of measuring the all-zeros state in the quantum circuit, $P(|0...0\rangle)$, equals the classically computed transformed cosine similarity S'_C. We validate this using quantum circuit simulations on the Qiskit Aer `qasm_simulator`, which emulates ideal quantum hardware. This approach allows us to focus solely on the theoretical framework's prediction, as validation on real quantum hardware would introduce a different research question centered on characterizing hardware noise and fidelity.

8.1. Embedding Model and Preprocessing

We use the Google `embedding-gemma-300m` model accessed through the `SentenceTransformer` library, which produces 768-dimensional embedding vectors. All embeddings are converted to 64-bit floating-point precision for numerical stability. As test inputs we use two semantically similar sentences: "Lions and tigers are both apex predators in their respective habitats" and "Tigers and lions, as top predators, occupy similar ecological niches." The raw embedding vectors a_{raw},$b_{raw} \in \mathbb{R}^{768}$ are L2-normalized so that $\|\mathbf{a}\| = \|\mathbf{b}\| = 1$ and the cosine similarity reduces to the dot product $S_C = \mathbf{a} \bullet \mathbf{b}$.

Since the quantum circuit requires the state dimension to be a power of 2, both vectors are zero-padded to $N_{padded} = 2^{10} = 1024$, requiring $n = 10$ qubits. Zero-padding preserves the cosine similarity because the appended components contribute nothing to the dot product or the norms. The vector $\mathbf{v}$ is constructed from the parameterization Eq. (3) as $v_i = (b_i + a_i)/\Delta$, where Δ is the minimum absolute value among the nonzero components of the raw embeddings. The theoretical prediction for the circuit output is S'_C (theoretical) = $(\mathbf{v} \cdot \mathbf{a}_{padded})^2/\|\mathbf{v}\|^2 = 0.8862$.

8.2. Results and Convergence

The quantum circuit was executed at various shot counts (N_{shots}). The measured probability $P(|0...0\rangle)$ consistently converged towards the theoretical S'_C value of 0.8862. For instance, at $N_{shots} = 8192$, the measured $P(|0...0\rangle)$ was 0.8842, with a deviation of 0.0020 from the theoretical value.

N_{shots}	$P(\|0\ldots0\rangle)$ measured	σ (expected)	$\|P - S'_C\|$
2048	0.8784	0.0070	0.0078
2048	0.8916	0.0070	0.0054
8192	0.8842	0.0035	0.0020

Table 1: Measured probability P(|0 . . . 0⟩) at different shot counts, compared to the theoretical prediction S'_C = 0.8862. The expected statistical uncertainty is $\sigma = \sqrt{S'_C(1 - S'_C)/N_{\text{shots}}}$.

The observed deviations are consistent with statistical shot noise, decreasing proportionally to $1/\sqrt{N_{\text{shots}}}$. This fundamental scaling for quantum measurements implies that arbitrary accuracy can be achieved by increasing the number of circuit executions. All measured values fell within their expected 95% confidence intervals, confirming that the differences are purely statistical.

8.3. Consistency and Validation

Numerical checks confirmed that zero-padding does not alter the transformed cosine similarity, ensuring that the quantum circuit's output corresponds to the original embedding vectors. The exact agreement between the classical cosine similarity, the Hamiltonian formula, and the quantum circuit measurements (within statistical fluctuations) validates the entire equivalence chain: $S'_C = \mathbf{a}^T\mathbf{H}'\mathbf{a} = a'^2_1 = |\tilde{\psi}_1|^2 = P(|0\ldots0\rangle)$. These results confirm the physical realizability of the framework in an ideal quantum environment.

8.4. Consolidated Consistency

The framework predicts that three independently computed quantities should agree: the classical cosine similarity S'_C (classical) = (1+**a•b**)/2, the Hamiltonian formula S'_C (Hamiltonian) = (**v•a**)^2/**‖v‖**^{2}, and the quantum measurement S'_C (quantum) = P(|0...0⟩). Table II presents the complete results.

Method	S'_C	**Uncertainty**
Classical (raw, $N = 768$)	0.8862	exact
Classical (padded, $N = 1024$)	0.8862	exact
Hamiltonian formula	0.8862	exact
Quantum circuit (2048 shots, run 1)	0.8784	±0.0070
Quantum circuit (2048 shots, run 2)	0.8916	±0.0070
Quantum circuit (8192 shots)	0.8842	±0.0035

Table 2: Consolidated results comparing classical computations and quantum circuit measurements for S'_C

The three classical computations agree exactly, confirming the internal consistency of the Hamiltonian formulation and the zero-padding invariance. The z-scores $z = |P - S'_C|/\sigma$ for the quantum measurements are 1.11, 0.77, and 0.57, all below 2, indicating that the deviations are well within the range expected from statistical fluctuations. The decreasing z-score with increasing shot count is consistent with convergence under the law of large numbers. These results validate each link in the equivalence chain $S'_C = \mathbf{a}^T\mathbf{H}'\mathbf{a} = a'^2_1 = |\tilde{\psi}_1|^2 = P(|0\ldots0\rangle)$.

9. Conclusions

This paper has established a structural isomorphism between LLM embedding spaces and quantum mechanical systems, providing a rigorous mathematical framework that connects classical cosine similarity to quantum measurement probability. The central result is that the transformed cosine similarity $S'_C = (1 + \mathbf{a} \bullet \mathbf{b})/2$ between two LLM embedding vectors admits four equivalent formulations connected by a commutative diagram (Figure 1)

$$S'_C = \mathbf{a}^T\mathbf{H}'\mathbf{a} = a'^2_1 = |\tilde{\psi}_1|^2 = P(|0\ldots0\rangle) \tag{75}$$

This equivalence chain demonstrates a formal identity between classical linear algebra, eigenspace decomposition, quantum amplitudes, and Born rule probabilities. The isomorphism is structural, not physical: it asserts that the mathematical structures governing LLM embeddings are formally identical to those of quantum mechanics, preserving algebraic, spectral, and observable properties. It does not imply that LLMs are physical quantum systems, nor does it suggest a computational speedup for pairwise similarity.

The framework rests on several key theoretical results. We provided a complete derivation of the LLM embedding classical system, showing how S'_C equals the squared projection onto an angle bisector, leading to a rank-1 Hamiltonian with a universal eigenstructure. The quantum system promotes this classical structure to complex, time-dependent state vectors governed by the Schrödinger equation.

We clarified that the time parameter is a gauge variable, not physical time, and that complex phases, while crucial for unitary evolution, cancel out in diagonal observables. The classical-to-quantum mapping is exact, preserving all observable quantities.

Symmetrically, the parameter $\hbar$ admits interpretation as the *semantic resolution* scale $\hbar_{\text{sem}} = \min_{i \neq j} \arccos(\mathbf{a}_i \bullet \mathbf{a}_j)$, the minimum angular separation between distinct token embeddings in the model's vocabulary. Only the ratio $\theta_s = t/\hbar$ enters observable predictions, so both parameters are gauge variables and the physically meaningful object is the semantic phase angle θ_s alone. This interpretation connects the abstract quantum parameter $\hbar$ to a directly computable property of any LLM embedding model, and yields the testable prediction that $\hbar_{\text{sem}}$ decreases monotonically with embedding dimension N.

The utility of this quantum mechanical language, despite its structural nature, is profound. It reframes cosine similarity as a measurement probability, connecting it to quantum measurement theory. It introduces complex phases, allowing for the definition of a semantic noise profile and, most significantly, reveals a local **U**(1) gauge symmetry [14]. This symmetry provides a mathematically precise model for contextual influence on embeddings, analogous to electromagnetism, with a conserved semantic charge [18]. Furthermore, the isomorphism enables direct implementation on quantum hardware, compressing N-dimensional embeddings onto $\lceil \log_2 N \rceil$ qubits, opening avenues for quantum algorithms operating natively on semantic representations.

The empirical content of this isomorphism is validated by quantum circuit simulations. The circuit, consisting of state preparation, unitary transformation, and measurement, produces $P(|0...0\rangle)$ which converges to the classically computed S'_C with deviations consistent with statistical shot noise. This confirms the physical realizability of the framework and provides a falsifiable prediction for any pair of LLM embedding vectors. The U(1) gauge structure offers further testable predictions regarding contextual perturbations.

In conclusion, this work provides a robust mathematical foundation for understanding LLM embedding spaces through the lens of quantum mechanics. It offers new conceptual tools and a richer framework for analyzing semantic representations, modeling contextual influences, and exploring quantum computing applications. The path forward involves leveraging these insights to develop new theoretical models and practical quantum algorithms for natural language processing.

References

1. Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *International Conference on Learning Representations*.
2. Pennington, J., Socher, R., & Manning, C. D. (2014, October). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1532-1543).
3. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems, 30*.
4. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1* (long and short papers) (pp. 4171-4186).
5. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems, 33,* 1877-1901.
6. Reimers, N., & Gurevych, I. (2019, November). Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)* (pp. 3982-3992).
7. Widdows, D., & Widdows, D. (2004). *Geometry and meaning* (Vol. 773). Stanford: CSLI publications.
8. Coecke, B., Sadrzadeh, M., & Clark, S. (2010). Mathematical foundations for a compositional distributional model of meaning. Lambek Festschrift Linguistic Analysis, 36(1).
9. Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum, 2,* 79.
10. Nielsen, M. A., & Chuang, I. L. (2010). *Quantum computation and quantum information* (10th anniversary ed.). Cambridge University Press.
11. Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing (STOC)* (pp. 212–219).
12. Kadowaki, T., & Nishimori, H. (1998). Quantum annealing in the transverse Ising model. *Physical Review E, 58*(5), 5355–5363.
13. Laine, T. A. (2025). Semantic Wave Functions: Exploring Meaning in Large Language Models Through Quantum Formalism. *OA J Applied Sci Technol, 3*(1), 01-22.
14. Laine, T. A. (2025). The Quantum LLM: Modeling Semantic Spaces with Quantum Principles. *OA J Applied Sci Technol, 3*(2), 01-13.
15. Laine, T. A. (2026). Quantum LLMs Using Quantum Computing to Analyze and Process Semantic Information. *OA J Applied Sci*

Technol, 4(1), 01-19.
16. Laine, T. A. (2026). Discrete semantic states and Hamiltonian dynamics in LLM embedding spaces. *OA Journal of Applied Science and Technology, 4*(1), 1–23.
17. Laine, T. A. (2026). Quantum computation of partition function similarity for large language models. *OA Journal of Applied Science and Technology, 4*(1), 1–11.
18. Laine, T. A. (2026). Quantum hierarchy for understanding LLM representations by modeling linear projections and nonlinear dynamics. *OA Journal of Applied Science and Technology, 4*(1), 1–43.
19. Laine, T. A. (2026). Quantum algorithms for large language models on noisy intermediate-scale quantum computers. *OA Journal of Applied Science and Technology, 4*(1), 1–13.